

CLAIMS

1. A multiprocessor computer system having a plurality of processors interconnected so that they can share memory, comprising:

a plurality of links, each link of said plurality of links connecting a processor to another processor;

a router box (RBOX) associated with each processor of said plurality of processors, said RBOX arranged to forward a message received on an input link of said plurality of links from a source processor to an outgoing link of said plurality of links toward a destination processor in response to data carried in said message;

a plurality of microprocessors, each of said microprocessors having a microprocessor memory associated therewith, a selected microprocessor of said plurality of microprocessors associated with at least one processor of said plurality of processors, said plurality of microprocessors arranged to control said plurality of processors, said control including applying electric power to a selected processor and removing electric power from said selected processor;

a data structure stored in microprocessor memory, said data structure storing a representation of the links connecting said processors and storing routes used by said RBOX in routing messages along said links, a copy of said data structure stored in microprocessor memory of each of said microprocessors; and,

a process to update said data structure in each said microprocessor memory in the event that a change occurs in a status of a component of said multiprocessor computer system.

2. The apparatus of claim 1 further comprising:

a process executing in said microprocessors for directing said microprocessors in formation of a partition of said processors, said partition having a selected number of said processors as members, said members capable of reading and writing a common memory within said partition, and other non-member processors excluded from reading and writing said common memory; and,

7 a second data structure for storing a representation of said partitions, and storing routes
8 through said links for transfer of messages between processors within a partition but not between
9 processors of different partitions.

1 3. The apparatus of claim 2 further comprising:

2 a management computer communicating through said local area network with said mi-
3 croprocessors, said management computer having an input device such as a keyboard and mouse
4 for entering commands to said plurality of microprocessors to modify said data base in order to
5 establish the processors belonging to a partition, said multiprocessor computer system supporting
6 a plurality of said partitions;

7 a process responsive to said data base, said process executing in said microprocessors to
8 establish that a processor in a partition can read and write memory associated with other proces-
9 sors of said partition, but cannot read and write memory associated with processors which are not
10 members of said partition.

1 4. The apparatus of claim 1 further comprising:

2 an input/output (IO) subsystem associated with selected processors of said plurality of
3 processors;

4 an IO microprocessor associated with each said IO subsystem, said IO microprocessor
5 communicating with said microprocessors through said local area network, each said IO micro-
6 processor having an IO microprocessor memory holding a copy of said database;

7 a process executing in said IO microprocessor, said process responsive to said database,
8 said process to apply power or remove power to said IO subsystem.

1 5. The apparatus of claim 1 further comprising:

2 a boot-up process executing in microprocessors of said plurality of microprocessors, said
3 boot-up process to start execution of said processors of said multiprocessor system, and said
4 multiprocessor system capable of continuing operation, after start of execution of said proces-
5 sors, to permit removal of a microprocessor from the multiprocessor system without interrupting
6 execution of said processors.

1 6. A method for operating a multiprocessor computer system, comprising:

2 connecting a plurality of processors so that they can share memory, a processor of said
3 plurality of processors connected to another processor by at least one link of a plurality of links;
4 associating a router box (RBOX) with each processor of said plurality of processors, said
5 RBOX arranged to forward a message received on an input link of said plurality of links from a
6 source processor to an outgoing link of said plurality of links toward a destination processor in
7 response to data carried in said message;

8 interconnecting a plurality of microprocessors, each of said microprocessors having a mi-
9 croprocessor memory associated therewith, a selected microprocessor of said plurality of micro-
10 processors associated with at least one processor of said plurality of processors, said plurality of
11 microprocessors arranged to control said plurality of processors, said control including applying
12 electric power to a selected processor and removing electric power from said selected processor;

13 storing a data structure in microprocessor memory, said data structure storing a repre-
14 sentation of the links connecting said processors and storing routes used by said RBOX in rout-
15 ing messages along said links, a copy of said data structure stored in microprocessor memory of
16 each of said microprocessors; and,

17 updating said data structure in each said microprocessor memory in the event that a
18 change occurs in a status of a component of said multiprocessor computer system.

1 7. The method of claim 6, further comprising:

2 executing a process in said microprocessors for directing said microprocessors in forma-
3 tion of a partition of said processors, said partition having a selected number of said processors
4 as members, said members capable of reading and writing a common memory within said parti-
5 tion, and other non-member processors excluded from reading and writing said common mem-
6 ory; and,

7 storing in a second data structure a representation of said partitions, and storing routes in
8 said second data structure through said links, said routes for transfer of messages between proc-
9 essors within a partition but not between processors of different partitions.

1 8. The method of claim 6, further comprising:

2 establishing communication between a management computer and said microprocessors
3 through said local area network, said management computer having an input device such as a
4 keyboard and mouse for entering commands to said plurality of microprocessors to modify said
5 data base in order to establish the processors belonging to a partition, said multiprocessor com-
6 puter system supporting a plurality of said partitions;

7 establishing, in response to said data base, that a processor in a partition can read and
8 write memory associated with other processors of said partition, but cannot read and write mem-
9 ory associated with processors which are not members of said partition.

1 9. The method of claim 6, further comprising:

2 associating an input/output (IO) subsystem with selected processors of said plurality of
3 processors;

4 associating an IO microprocessor with each said IO subsystem, said IO microprocessor
5 communicating with said microprocessors through said local area network, each said IO micro-
6 processor having an IO microprocessor memory holding a copy of said database;

7 executing a process in said IO microprocessor, said process responsive to said database,
8 said process to apply power or remove power to said IO subsystem.

1 10. The method as in claim 6, further comprising:

2 executing a boot-up process in microprocessors of said plurality of microprocessors, said
3 boot-up process to start execution of said processors of said multiprocessor system, and said
4 multiprocessor system capable of continuing operation, after start of execution of said proces-
5 sors, to permit removal of a microprocessor from the multiprocessor system without interrupting
6 execution of said processors.